



Clear and Present Danger of IPv6

- IPv6/IPv4 fallback and DNS queries-

Tomohiro Fujisaki
fujisaki@nttv6.net

NTT Information Sharing Platform Labs.

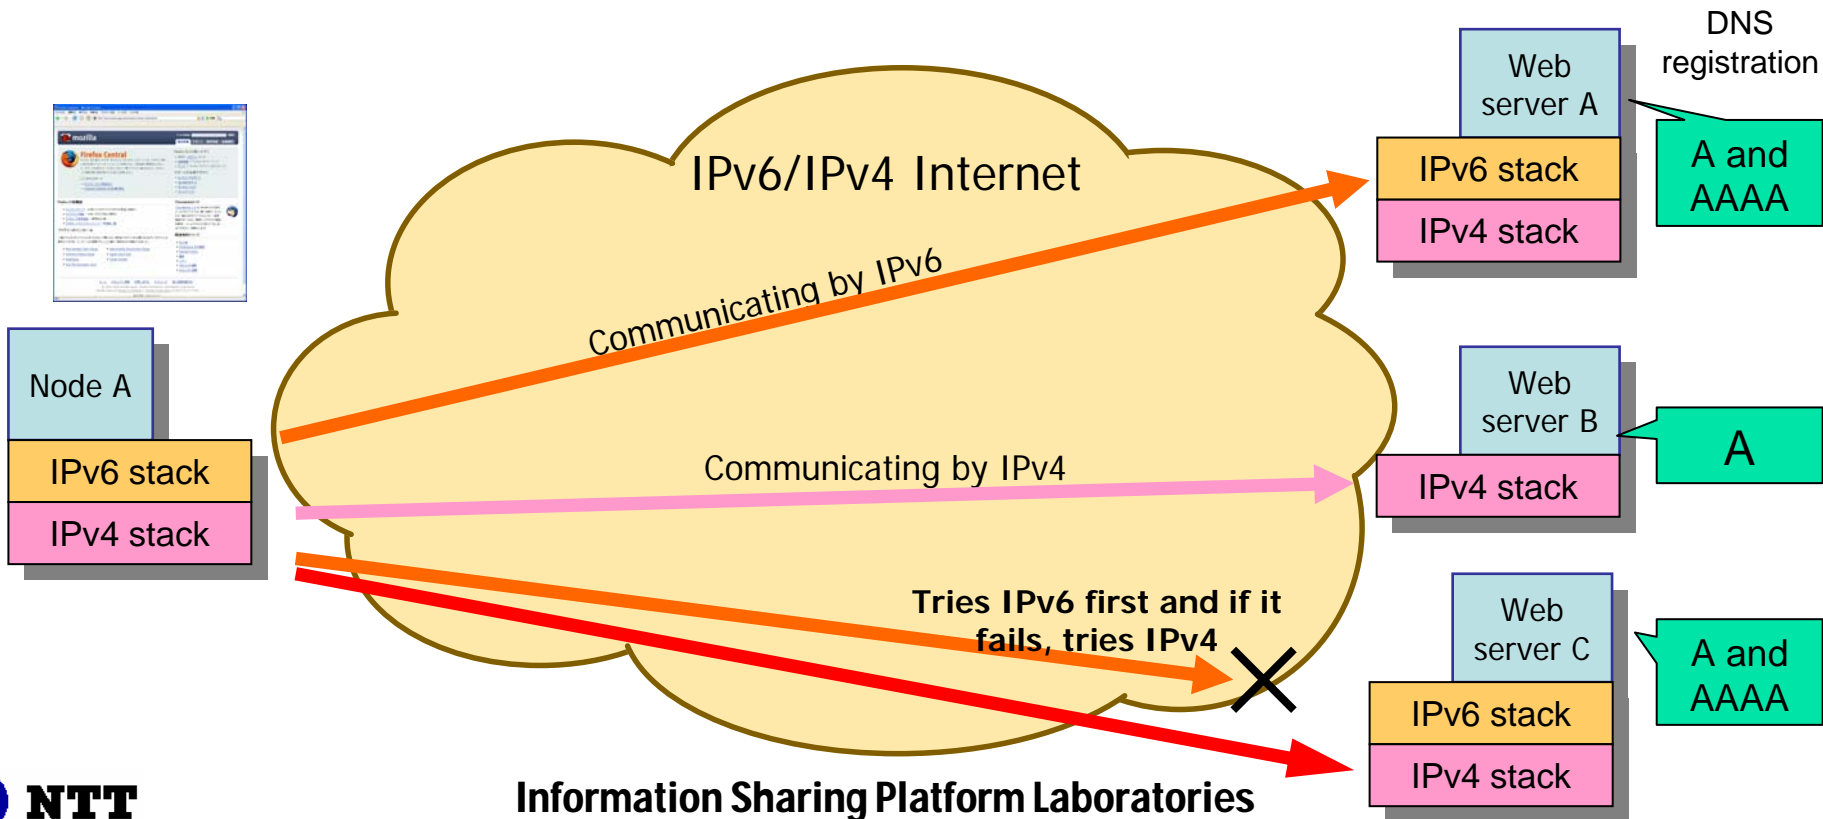


Network problems associated with IPv6

- IPv6 is widely used
 - Many operating systems and routers support IPv6 without any additional software.
 - Anyone can use IPv6
- IPv6/IPv4 fallback problem
 - Described in section 4 at <http://www.v6fix.net/docs/wide-draft-v6fix.en>
- DNS cache server problem

What is 'IPv6/IPv4 fallback'?

- Many current IPv6/IPv4 dual-stack operating systems start their communication using IPv6.
 - If IPv6 communication fails, they try to communicate using IPv4.
 - Dual-stack node application behavior is IPv6-to-IPv4 fallback.




IPv6/IPv4 fallback problem

Problem: Fallback sometimes takes a long time.


- Problem is especially significant in TCP-based applications.
 - They initially need to establish communication channels
 - If destination node has multiple IP addresses, application tries them sequentially until TCP communication channel is established.

User's
view

User has to wait until communication is finished
until web page is displayed completely.



Over 20 seconds...



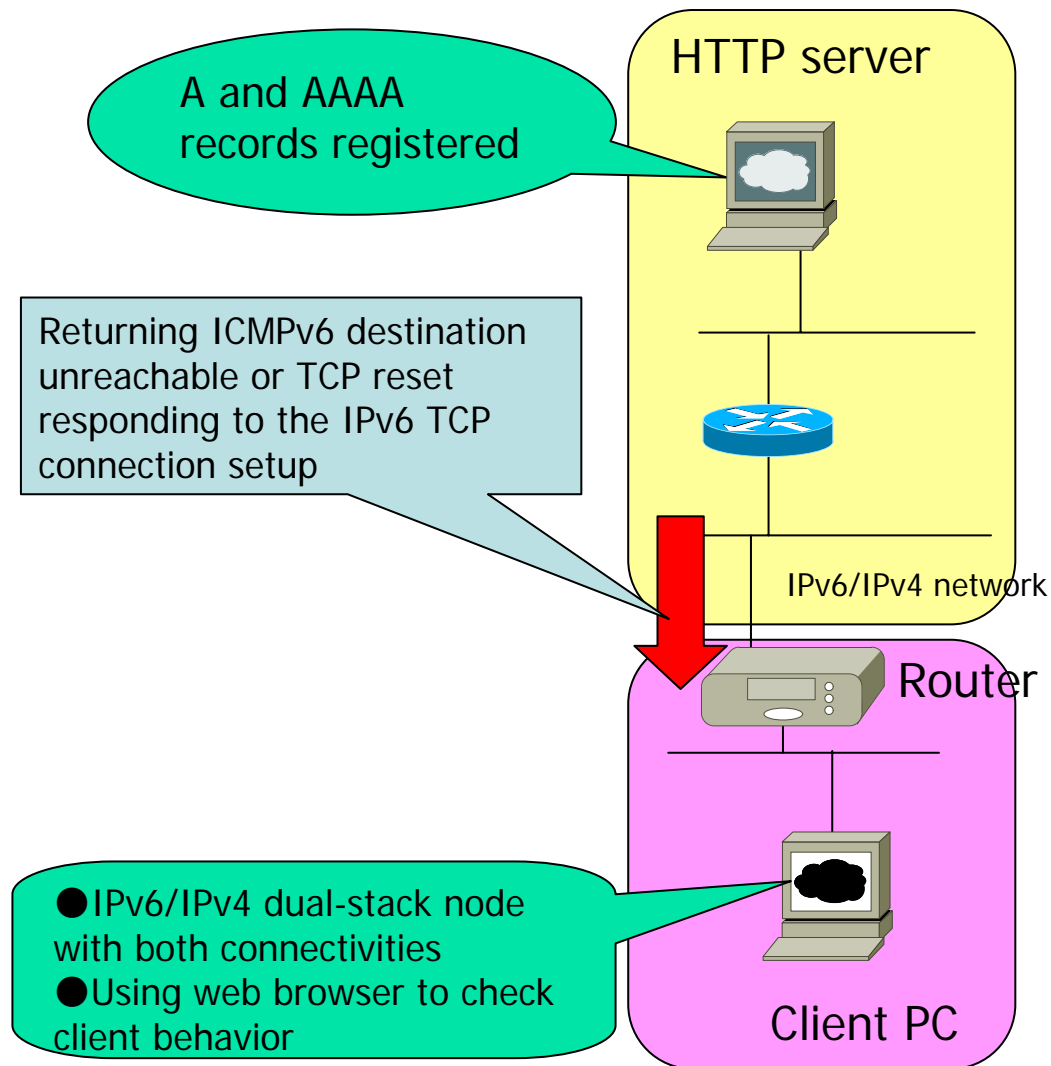


What causes IPv6/IPv4 fallback?

- Poor IPv6 connectivity.
 - Using unmanaged transition mechanism
 - Destination server has AAAA records, but no IPv6 connectivity
 - No IPv6 global connectivity (e.g., VPNs)
- TCP stack judges 'failure of connection' by:
 - Timeout of connection
 - Error reports from network nodes
 - ICMP error
 - Mainly 'ICMP destination unreachable' message
 - TCP RST

Fallback experimental setup

- Measuring time required to fall back from IPv6 to IPv4 at client PC when:
 - No response from network
 - ICMPv6 destination unreachable returned
 - no route to dest
 - administratively prohibited
 - address unreachable
 - port unreachable
 - TCP RST returned



Results of experiments 1/3

unit: seconds

Operating systems	Timeout	ICMPv6 error (Type=1: Destination unreachable)				TCP RST
		Code=0 No Route	Code=3 Addr Unreach	Code=1 Admin Prohibit	Code=4 Port Unreach	
WinXP (SP2) IE 6.x	22.87	22.91	22.92	22.95	22.93	1.06
WinXP (SP2) Firefox 1.5	20.99	20.96	20.96	20.98	No fallback	0.97
Longhorn (Vista) (IE 7.0.112)	21.00	21.00	21.00	21.00	21.01	1.00

⌘ Time between first IPv6 TCP SYN packet and IPv4 TCP SYN packet immediately after fallback occurs.

Results of experiments 2/3

unit: seconds

Operating systems	Timeout	ICMPv6 Error (Type=1: Destination unreachable)				TCP RST
		Code=0 No Route	Code=3 Addr Unreach	Code=1 Admin Prohibit	Code=4 Port Unreach	
Fedora-C3(kernel:2.6.9) Mozilla1.7.8	188.98	1.65	1.91	No fallback	0.23	0.00
Fedora-C3(kernel:2.6.11) FireFox1.5	186.41	0.01	0.00	0.00	0.01	0.01
FreeBSD-R5.4 Mozilla1.7.8	75.01	12.60	No fallback	12.60	12.60	1.38
FreeBSD-R5.4 Mozilla1.7.12	75.01	12.61	No fallback	12.60	12.60	0.01
FreeBSD-R5.4 (with 20050620 kame) Mozilla1.7.8	75.00	1.78	No fallback	12.61	12.61	0.90
FreeBSD-R5.4 (with 20060116 kame) Mozilla1.7.12	75.00	0.02	No fallback	12.62	12.60	0.03
MacOS X Tiger Safari2.0.3	149.58	11.89	11.85	11.55	11.99	0.08
Solaris10 Mozilla1.7.12	224.29	224.69	224.69	224.69	0.01	1.05

Results of experiments 3/3


unit: seconds

Operating systems	Timeout	ICMPv6 Error (Type=1: Destination Unreachable)				TCP RST
		Code=0 No Route	Code=3 Addr Unreach	Code=1 Admin Prohibit	Code=4 Port Unreach	
FreeBSD-R5.4 FireFox1.5	75.00	12.60	No fallback	12.60	12.61	0.01
FreeBSD-R5.4 (with 20060116 kame) FireFox1.5	75.00	0.01	No fallback	12.60	12.60	0.02
MacOS X Tiger FireFox1.5	74.74	11.80	No fallback	11.72	11.63	0.01
Solaris10 FireFox1.5	224.69	224.69	224.69	224.69	0.02	0.02



TCP stack behavior when ICMP errors are returned

- When node receives ICMP error packets, TCP stack behavior is defined in RFC 1122 (for IPv4 ICMP only).
 - When node receives an ICMP hard error, TCP aborts connection immediately. When node receives ICMP soft error, TCP must not abort connection.
- The “draft-gont-tcpm-tcp-soft-errors-*.txt” proposes IPv6 version of ICMPv6 soft error handling.



Currently, ICMPv6 destination unreachable error handling is different in each OS.



Summary of IPv6/IPv4 fallback problem

- In IPv6-and-IPv4 enabled network, quality of each network affects the user's nodes.
 - If quality (e.g. reachability) of the network is clearly bad, administrators should control communication by:
 - Changing protocol preference using RFC3484 mechanism
 - Returning TCP RST to minimize fallback time




DNS cache server issue: Increase in Number of DNS AAAA Queries

This topic was presented by T. Toyono
and K. Toyama at NANOG 36

<http://www.nanog.org/mtg-0602/ishibashi.html>

Client's DNS query problem

- Focus on increase in number of queries between user resolvers and DNS cache servers caused by:
 - 1. IPv6 support
 - Same number of AAAA queries as A queries
 - 2. Domain name completion
 - By operating system
 - By applications





1. IPv6 and OS resolver

- IPv6-enabled OSs send AAAA queries for every name resolution
- BSD UNIX / Windows
 - Send A and AAAA queries for every name resolution
 - Almost all applications do not specify “DNS Query Type”, so OS sends both.
 - Even if response to AAAA query is “not exist such a domain name” (NXDomain), OS tries to send A query.



2. Domain name completion

- When name resolution fails, OS and APP automatically attempt to resolve domain with prefix/suffix completion.
 - e.g., when name resolution of “host” fails...
→ host.com → host.org → host.net ...
- OS using these domains to complete:
 - FreeBSD: specified by “search” in /etc/resolv.conf and distributed via DHCP
 - Windows: configured in control panel and distributed via DHCP
- Applications:
 - Mozilla: tries to resolve domain name by adding “www. ” domain prefix
 - IE: searches domain using MSN search and then retries name resolutions for domains by adding .com → .org → .net → .edu

Combinations in FreeBSD

- Combinations of AAAA queries and domain name completions are different depending on OS.
- FreeBSD
 - Tries domain name completions for A and AAAA

(E.g.) User Query: noexist-example.com

A noexist-example.com

AAAA noexist-example.com

A noexist-example.com.com

AAAA noexist-example.com.com

A noexist-example.com.net

AAAA noexist-example.com.net

If IPv4 address is resolved,
stop here.



Combinations in Windows

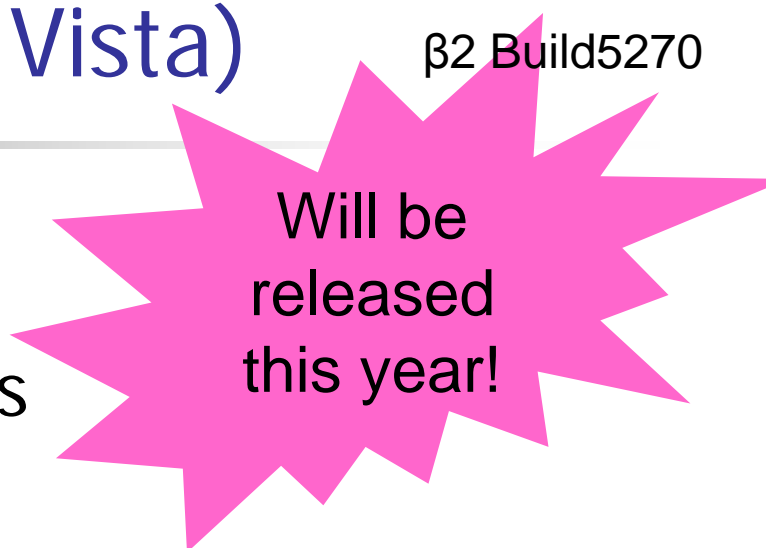
- Combinations of AAAA queries and domain name completions are different depending on OS.
- Windows
 - Tries AAAA queries for all domain completions, and then A queries with domain completions.

(E.g.) User Query: noexist-example.com
AAAA noexist-example.com
AAAA noexist-example.com.com
AAAA noexist-example.com.net
A noexist-example.com
A noexist-example.com.com
A noexist-example.com.net



Longhorn (Windows Vista)

β2 Build5270



Will be
released
this year!

- **Default IPv6 enabled**
 - Always tries AAAA queries
- OS/Application domain name completion
 - Behavior of OS resolver is same as that of Windows XP
 - OS and applications make unnecessary suffix/prefix completions for domain names.

IPv6 enabled + domain name completion common in IE7

```
AAAA noexist.nttv6
```

```
response, No such name
```

```
AAAA noexist.nttv6.suffix.os.nttv6.org
```

```
response, No such name
```

```
AAAA noexist.nttv6.suffix.interface.nttv6.net
```

```
response, No such name
```

```
AAAA noexist.nttv6.os.nttv6.org
```

```
response, No such name
```

```
AAAA noexist.nttv6.nttv6.org
```

```
response, No such name
```

```
A noexist.nttv6
```

```
response, No such name
```

```
A noexist.nttv6.suffix.os.nttv6.org
```

```
response, No such name
```

```
A noexist.nttv6.suffix.interface.nttv6.net
```

```
response, No such name
```

```
A noexist.nttv6.os.nttv6.org
```

```
response, No such name
```

```
A noexist.nttv6.nttv6.org
```

```
response, No such name
```

```
AAAA search.msn.com
```

```
response CNAME search.msn.com.nsatc.net CNAME search.msn.com.edge.suite.net CNAME a134-p.akamai.net
```

```
A search.msn.com
```

```
response CNAME search.msn.com.nsatc.net CNAME search.msn.com
```

Only 1 user click...

IPv6 address

X

domain have completions

IPv4 address

X

domain name completions

MSN search

using AAAA and A

A 61

- 1 user click → 12 DNS queries

Worst case...

```
AAAA noexist.nttv6
AAAA noexist.nttv6.suffix.os.nttv6.org
AAAA noexist.nttv6.suffix.interface.nttv6.net
AAAA noexist.nttv6.os.nttv6.org
AAAA noexist.nttv6.nttv6.org
A noexist.nttv6
A noexist.nttv6.suffix.os.nttv6.org
A noexist.nttv6.suffix.interface.nttv6.net
A noexist.nttv6.os.nttv6.org
A noexist.nttv6.nttv6.org
AAAA auto.search.msn.com
A auto.search.msn.com
AAAA sea.search.msn.co.jp
AAAA www.noexist.nttv6.co.jp
AAAA www.noexist.nttv6.co.jp.suffix.os.nttv6.org
AAAA www.noexist.nttv6.co.jp.suffix.interface.nttv6.net
AAAA www.noexist.nttv6.co.jp.os.nttv6.org
AAAA www.noexist.nttv6.co.jp.nttv6.org
A www.noexist.nttv6.co.jp
A www.noexist.nttv6.co.jp.suffix.os.nttv6.org
A www.noexist.nttv6.co.jp.suffix.interface.nttv6.net
A www.noexist.nttv6.co.jp.os.nttv6.org
A www.noexist.nttv6.co.jp.nttv6.org
AAAA www.noexist.nttv6.org
AAAA www.noexist.nttv6.org.suffix.interface.nttv6.net
A www.noexist.nttv6.org
A www.noexist.nttv6.org.suffix.interface.nttv6.net
AAAA www.noexist.nttv6.net
AAAA www.noexist.nttv6.net.suffix.os.nttv6.org
www.noexist.nttv6.net.suffix.os.nttv6.org
www.noexist.nttv6.net.suffix.interface.nttv6.net
www.noexist.nttv6.net.suffix.os.nttv6.org
www.noexist.nttv6.net.suffix.interface.nttv6.net
www.noexist.nttv6.net.suffix.os.nttv6.org
www.noexist.nttv6.net.suffix.interface.nttv6.net
www.noexist.nttv6.net.suffix.os.nttv6.org
www.noexist.nttv6.net.suffix.interface.nttv6.net
www.noexist.nttv6.edu.os.nttv6.org
www.noexist.nttv6.edu.nttv6.org
AAAA sea.search.msn.co.jp
AAAA sea.search.msn.co.jp
```

One user click



40 queries...

OS domain name completion

IE tried MSN search

IE added “.com”
and OS domain completion

IE added “.net”
and OS domain completion

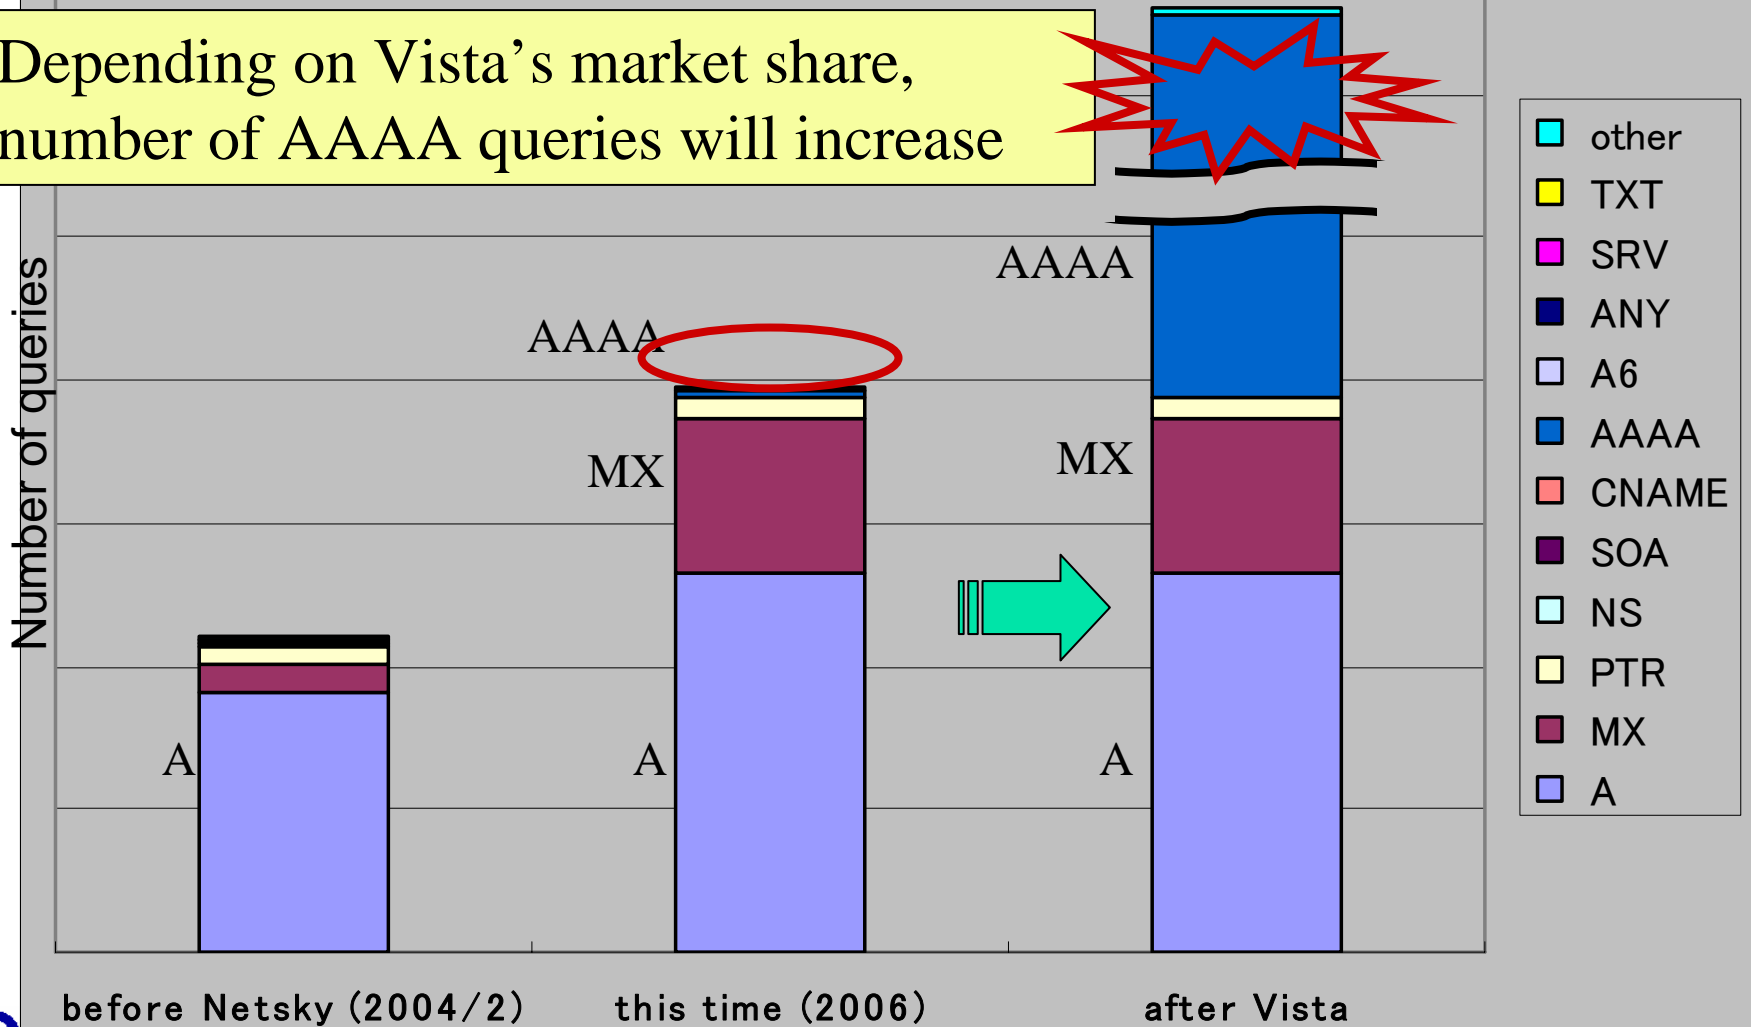
IE added “.org”
and OS domain completion

IE added “.edu”
and OS domain completion

IE tried MSN search

Expected increase in number of user queries

Depending on Vista's market share, number of AAAA queries will increase





Summary of DNS cache server issue

- IPv6-enabled operating systems double number of user DNS queries
 - Increases number of queries in domain name completions and domain search sequences for AAAA and A queries
- Suggestions to operators
 - Cache servers should be prepared for the increases.
 - DNS response time has a large impact on QoS to end users e.g., stopping domain name distribution to users by DHCP or PPP